arXiv:1404.0329v1 [q-bio.QM] 1 Apr 2014

Toward computational cumulative biology by combining models of biological datasets

Ali Faisal [*], **Jaakko Peltonen** [*], **Elisabeth Georgii** [*], **Johan Rung** [†] **and Samuel Kaski** [*] [‡]

[*]Helsinki Institute for Information Technology HIIT, Department of Information and Computer Science, PO Box 15400, Aalto University, 00076 Aalto, Finland,[†]European Molecular Biology Laboratory, European Bioinformatics Institute (EMBL-EBI), Wellcome Trust Genome Campus Hinxton, CB10 1SD, UK, and [‡]Helsinki Institute for Information Technology HIIT, Department of Computer Science, PO Box 68, 00014 University of Helsinki, Finland

A main challenge of data-driven sciences is how to make maximal use of the progressively expanding databases of experimental datasets in order to keep research cumulative. We introduce the idea of a modeling-based dataset retrieval engine designed for relating a researcher's experimental dataset to earlier work in the field. The search is (i) data-driven to enable new findings, going beyond the state of the art of keyword searches in annotations, (ii) modeling-driven, to both include biological knowledge and insights learned from data, and (iii) scalable, as it is accomplished without building one unified grand model of all data. Assuming each dataset has been modeled beforehand, by the researchers or by database managers, we apply a rapidly computable and optimizable combination model to decompose a new dataset into contributions from earlier relevant models. By using the data-driven decomposition we identify a network of interrelated datasets from a large annotated human gene expression atlas. While tissue type and disease were major driving forces for determining relevant datasets, the found relationships were richer and the model-based search was more accurate than keyword search; it moreover recovered biologically meaningful relationships that are not straightforwardly visible from annotations, for instance, between cells in different developmental stages such as thymocytes and T-cells. Data-driven links and citations matched to a large extent; the data-driven links even uncovered corrections to the publication data, as two of the most linked datasets were not highly cited and turned out to have wrong publication entries in the database.

bioinformatics | generative models | information retrieval | machine learning

Significance

Measurement data sets of molecular biology and other experimental sciences are being collected comprehensively to openly accessible databases. We demonstrate that it is now possible to relate new results to earlier science by searching with the actual data, instead of only in the textual annotations which are restricted to known findings and are the current state of the art. In a gene expression database, the data-driven relationships between datasets matched well with citations between the corresponding research papers, and even found mistakes in the database.

Molecular biology, historically driven by the pursuit of experimentally characterizing each component of the living cell, has been transformed into a data-driven science [1, 2, 3, 4, 5, 6] with just as much importance given to the computational and statistical analysis as to experimental design and assay technology. This has brought to the fore new computational challenges, such as the processing of massive new sequencing data, and new statistical challenges arising from the problem of having relatively few (n) samples characterized for relatively many (p) variables—the "large p, small n" problem. High throughput technologies often are developed to assay many parallel variables for a single sample in a run, rather than many parallel samples for a single variable, whereas the statistical power to infer properties of biological conditions increases with larger sample sizes. For cost reasons, most labs are restricted to generating datasets with the statistical power to detect only the strongest effects. In combination with the penalties of multiple hypothesis testing, the limitations of "large p, small n" datasets are obvious. It is therefore not surprising that much work has been devoted to address this problem.

Some of the most successful methods rely on increasing the effective number of samples by combining with data from other, similarly designed, experiments, in a large meta-analysis [7]. Unfortunately, this is not straightforward either. Although public data repositories, such as the ones at NCBI in the United States and the EBI in Europe, serve the research community with ever-growing amounts of experimental data, they largely rely on annotation and meta-data provided by the submitter. Database curators and semantic tools such as ontologies provide some help in harmonizing and standardizing the annotation, but the user who wants to find datasets that are combinable with her own, most often must resort to searches in free text or in controlled vocabularies which would need much downstream curation and data analysis before any meta-analysis can be done [8].

Ideally, we would like to let the data speak for themselves. Instead of searching for datasets that have been described similarly, which may not correspond to a statistical similarity in the datasets themselves, we would like to conduct that search in a data-driven way, using the dataset itself as the query, or a statistical (rather than a semantic) description of it. This is implicitly done for example in multi-task learning, a method from the machine learning field [9, 10], where several related estimation tasks are pursued together, assuming shared properties across tasks. Multi-task learning is a form of global analysis, which builds a single unified model of the datasets. But as the number of datasets keeps increasing and the amount of quantitative biological knowledge keeps accumulating, the complexity of building an accurate unified model becomes increasingly prohibitive.

Addressing the "large p, small n" problem requires taking into account both the uncertainty in the data and the existing biological knowledge. We now consider the hypothesized scenario where future researchers increasingly develop hypotheses in terms of (probabilistic) models of their data. Although far from realistic today, a similar trend exists for sequence motif data, which are often published as Hidden Markov models, for instance in the Pfam database [11].

In this paper we report on a feasibility study towards the scenario where a large number of experiments have been modeled beforehand, potentially by the researcher generating the data or the database storing the model together with the data. We ask *what could be done with these models towards cumulatively building knowledge from data in molecular biology*. Speaking about models generally and assuming the many practical issues can be solved technically, we arrive at our answer: *a modeling-driven dataset retrieval engine*, which a researcher can use for positioning her own measurement data into the context of the earlier biology. The engine will point out relationships between experiments in the form of the retrieval results, which is a naturally understandable interface. The retrieval will be based on data instead of the state of the art of using keywords and ontologies, which will make unexpected and previously unknown findings possible. The retrieval will use the models of the datasets which, by our assumption above, incorporate what the researchers producing the data thought was important, but the retrieval will be designed to be more scalable than building one unified grand model of all data. This also implies that the way the models are utilized needs to be approximate. Compared to existing data-driven retrieval methods

[5, 3], whole datasets, incorporating the experimental designs, will be matched instead of individual observations. The remaining question is how to design the retrieval so that it both reveals the interesting and important relationships and is fast to compute.

The model we present is a first step towards this goal. We assume a new dataset can be explained by a combination of the models for the earlier datasets and a novelty term. This is a mixture modeling or regression task, in which the weights can be computed rapidly; the resulting method scales well to large numbers of datasets, and the speed of the mixture modeling does not depend on the sizes of the earlier datasets. The largest weights in the mixture model point at the most relevant earlier datasets. The method is applicable to several types of measurement datasets, assuming suitable models exist. Unlike traditional mixture modeling we do not limit the form of the mixture components, thus we bring in the knowledge built into the stored models of each dataset. We apply this approach to a large set of experiments from EBI's ArrayExpress gene expression database [12], treating each experiment in turn as a new dataset, queried against all earlier datasets. Under our assumptions the retrieval results can be interpreted as studies that the authors of the study generating the query set could have cited, and we show that the actual citations overlap with the retrieval results. The discovered links between datasets additionally enable forming a "hall of fame" of gene expression studies, containing the studies that would have been influential assuming the retrieval system would have existed. The links in the "hall of fame" verify and complement the citation links: in our study they revealed corrections to the citation data, as two frequently retrieved studies were not highly cited and turned out to have erroneous publication entries in the database. We provide an online resource for exploring and searching this "hall of fame": `http://research.ics.aalto.fi/mi/setretrieval`.

Earlier work on relating datasets has provided partial solutions along this line, with the major limitation of being restricted to pairwise dataset comparisons, in contrast to the proposed approach of decomposing a dataset into contributions from a set of earlier datasets. Russ and Futschik [13] represented each dataset by pairwise correlations of genes, and used them to compute dataset similarities. This dataset representation is ill-suited for typical functional genomics experiments as a large number of samples is required to sensibly estimate gene correlation matrices. In addition, it makes the dataset comparison computationally expensive, as the representation is bulkier than the original dataset. In other works specific case-control designs [14] or known biological processes [15] are assumed; we generalize to decompositions over arbitrary models.

Combination of stored models for dataset retrieval

Our goal is to infer data-driven relationships between a new "query" dataset q and earlier datasets. The query is a dataset of N_q samples $\{x_i^q\}_{i=1}^{N_q}$; in the ArrayExpress study the samples are gene expression profiles, with the element x_{ij}^q being expression of the gene set j in the sample i of the query q, but the setup is general and applicable to other experimental data as well. Assume further a dataset repository of N_S earlier datasets, and assume that each dataset s_j, $j = 1, \ldots, N_S$, has already been modeled with a model denoted by M^{s_j}, later called a base model. The base models are assumed to be probabilistic generative models, *i.e.*, principled data descriptions capturing prior knowledge and data-driven discoveries under specific distributional assumptions. Base models for different datasets may come from different model families, as chosen by the researchers who authored each dataset. In this paper we use two types of base models, which are discrete variants of principal component analysis (*Results*), but any probabilistic generative models can be applied.

Assume tentatively that the dataset repository contains a library of "base experiments", carefully selected to induce all important known biological effects with suitable design factors. In the special example case of metagenomics with known constituent organisms, an obvious set of base experiments would be the set of genomes of those organisms [16]. A new experiment could then be expressed as a combination of the base experiments, and potential novel effects. More generally, for instance in a broad gene expression atlas, it would be hard if not impossible to settle on a clean, well-defined and up-to-date base set of experiments to correspond to each known effect, so we choose to *use the comprehensive collection of experiments in the current databases as the base experiments*. The problem setting then changes, from searching for a unique explanation of the new experiment, to the down-to-earth and realistic task of data-driven retrieval of a set of relevant earlier experiments, relevant in the sense of having induced one or more of the known or as-of-yet unknown biological effects.

We combine the earlier datasets by a method that is probabilistic but simple and fast. We build a *combination model* for the query dataset as a mixture model of base distributions $p(x|M^{s_j})$, which have been estimated beforehand. In our scenario, generative models M^{s_j} are available in the repository along with datasets s_j; note that the M^{s_j} need not all have the same form. In the mixture model parameterized by $\Theta^q = \{\theta_j^q\}_{j=1}^{N_S+1}$, the likelihood of observing the query is

$$p(\{x_i^q\}_{i=1}^{N_q}; \Theta^q) = \prod_{i=1}^{N_q} \Big[\Big(\sum_{j=1}^{N_S} \theta_j^q p(x_i^q|M^{s_j}) \Big) + \theta_{N_S+1}^q p(x_i^q|\psi) \Big] \quad \textbf{[1]}$$

where θ_j^q is the mixture proportion or *weight* of the jth base distribution (model of dataset s_j) and $\theta_{N_S+1}^q$ is the weight for the novelty term. The novelty is modeled by a "background model" ψ, a broad nonspecific distribution covering overall gene-set activity across the whole dataset repository. All weights are non-negative and $\sum_{j=1}^{N_S+1} \theta_j^q = 1$. In essence, this representation assumes that biological activity in the query dataset can be approximately explained as a combination of earlier datasets and a novelty term.

The remaining task is to infer the combination model Θ^q for each query q given the known models M^{s_j} of datasets in the repository. We infer a maximum a posteriori (MAP) estimate of the weights $\Theta^q = \{\theta_j^q\}_{j=1}^{N_S+1}$. Alternatively we could sample over the posterior, but MAP inference already yielded good results. We optimize the combination weights to maximize their (log) posterior probability

$$\log p(\{\theta_j^q\}|\{x_i^q\}, \{M^{s_j}\}) \propto \log p(\{x_i^q\}|\{M^{s_j}\}, \{\theta_j^q\}) + \log p(\{\theta_j^q\})$$

$$\propto \sum_i \log \Big[\Big(\sum_{j=1}^{N_S} \theta_j^q p(x_i^q|M^{s_j}) \Big) + \theta_{N_S+1}^q p(x_i^q|\psi) \Big] - \lambda \sum_{j=1}^{N_S+1} \theta_j^{q2} \quad \textbf{[2]}$$

where $p(\{\theta_j^q\}) = \mathcal{N}(0, \lambda^{-1} I)$ is a naturally non-sparse L_2 prior for the weights with a regularization term λ. The cost function **[2]** is strictly concave (*SI Text*), and standard constrained convex optimization techniques can be used to find the optimized weights. Algorithmic details for the Frank-Wolfe algorithm and a proof of convergence are provided in *SI Text*. After computing the MAP estimate, we rank the datasets for retrieval according to decreasing combination weights.

This modeling-driven approach has several advantages: 1) the approximations become more accurate as more datasets are submitted to the repository, increasing naturally the number of base distributions; 2) it is fast, since only the models of the datasets are needed, not the large datasets themselves; 3) any model types can be included, as long as likelihoods of an observed sample can be computed, hence all expert knowledge built into the models in the repository can be used; 4) relevant datasets are not assumed to be "similar" to the query in any naïve sense, they only need to explain a part of the query set; 5) the relevance scores of datasets have a natural quantitative meaning as weights in the probabilistic combination model.

Scalability. As the size of repositories such as ArrayExpress doubles every two years or even faster [17], fast computation with respect

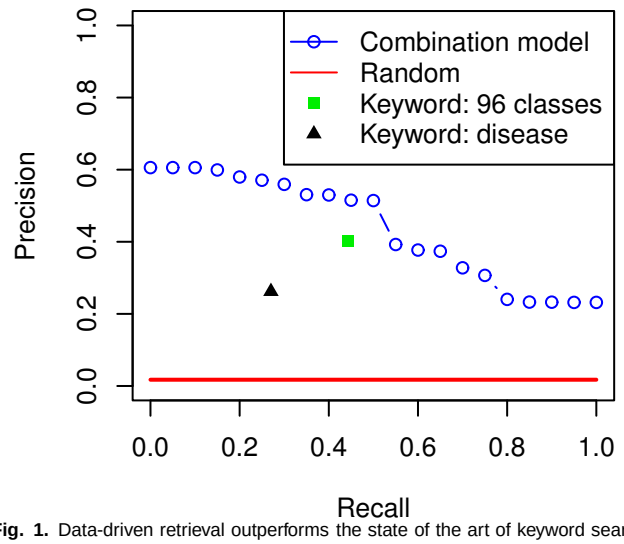

Fig. 1. Data-driven retrieval outperforms the state of the art of keyword search on the human gene expression atlas [12]. Blue: Traditional precision-recall curve where progressively more datasets are retrieved from left to right. All experiments sharing one or more of the 96 biological categories of the atlas were considered relevant. In keyword retrieval, either the category names ("Keyword: 96 classes") or the disease annotations ("Keyword: disease") were used as keywords. All datasets having at least ten samples were used as query datasets, and the curves are averages over all queries.

to the number N_S of background datasets is crucial for future-proof search methods. Already the first method above has a fast linear computation time in N_S (*SI Text*), and an approximate variant can be run in sublinear time. For that, the model combination will be optimized only over the k background datasets most similar to the query, which can be found in time $O(N_S^{1/(1+\varepsilon)})$ where $\varepsilon \geq 0$ is an approximation parameter [18], by suitable hashing functions.

Results

Data-driven retrieval of experiments is more accurate than standard keyword search. We benchmarked the combination model against state-of-the-art dataset retrieval by keyword search, in the scenario where a user queries with a new dataset against a database of earlier released datasets represented by models. The data were from a large human gene expression atlas [12], containing 206 public datasets with 5372 samples in total that have been systematically annotated and consistently normalized. To make use of prior biological knowledge, we preprocessed the data by gene set enrichment analysis [19], representing each sample by an integer vector telling for each gene set the number of leading edge active genes [20] (*Methods*). As base models we used two model types previously applied in gene expression analysis [6, 20, 21, 3]: a discrete principal component analysis method called Latent Dirichlet Allocation [22, 23], and a simpler variant called mixture of unigrams [24] (*SI Text*). Of the two types, for each dataset we chose the model yielding the larger predictive likelihood (*SI Text*). For each query (q), the earlier datasets (s_j) were ranked in descending order of the combination proportion (θ_j^q; estimated from Eq. **[2]**). That is, base models which explained a larger proportion of the gene set activity in the query were ranked higher. The approach yields good retrieval: the retrieval result was consistently better than with keyword searches applied to the titles and textual descriptions of the datasets (Fig. 1), which is a standard approach for dataset retrieval from repositories [25].

We checked that the result is not only due to laboratory effects by discarding, in a follow-up study, all retrieved results from the same laboratory. The mean average precision decreased slightly (from 0.44 to 0.42; precision-recall curve in Fig. S2) but still supports the same conclusion.

Network of computationally recommended dataset connections reveals biological relationships. When each dataset in turn is used as a query, the estimated combination weights form a "relevance network" between datasets (Fig. 2 left), where each dataset is linked to the relevant earlier datasets (for details, see *Methods*; a larger figure in Fig. S5; and an interactive searchable version in the online resource). The network structure is dominated but not fully explained by the tissue type. Normal and neoplastic solid tissues (cluster 1) are clearly separate from cell lines (cluster 2) and from hematopoietic tissue (cluster 4); the same main clusters were observed in [12]. Note that the model has not seen the tissue types but has found them from the data. In closer inspection of the clusters, some finer structure is evident. The muscle and heart datasets (gray) form an interconnected subnetwork in the left edge of the image: nodes near the bottom of the image (downstream) are explained by earlier (upstream) nodes, which in turn are explained by nodes even further upstream. As another example, in cluster 4 myeloma and leukemia datasets are concentrated on the left side of the cluster, whereas the right side mostly contains normal or infected mononuclear cells.

There is a substantial number of links both across clusters and across tissue categories. Among the top thirty cross-category links, 25 involve heterogeneous datasets containing samples from diverse tissue origins. The strongest link connects GSE6365, a study on multile myeloma, with GSE2113, a larger study from the same lab which largely includes the GSE6365 samples. The dataset E-MEXP-66 is a hub connected to all the clusters and to nodes in its own cluster having different tissue labels. It contains samples studying Kaposi sarcoma, and includes control samples from skin endothelial cells from blood vessels and the lymph system. Blood vessels and cells belonging to the lymph system are expected to be present in almost any solid tissue biopsy as well as in samples based on blood samples. The strongest link between two homogeneous datasets of different tissue types connects GSE3307 (which compares skeletal muscle samples from healthy individuals with 12 groups of patients affected by various muscle diseases) to GSE5392, which measures transcriptome profiles of normal brain and brain bipolar disorder. Interestingly, shortening of telomeres has been associated both with bipolar disorder [28] and muscular disorder [26]. Treatment of bipolar disorder has been found to also slow down the onset of skeletal muscle disorder [27].

Next we investigated "outlier" datasets where the tissue type does not match the main tissue types of a cluster, implying that they might reveal commonalities between cellular conditions across tissues. Cluster 1 contained three outlier datasets: two hematopoietic datasets and one cell line dataset. The two hematopoietic outlier datasets are studies related to macrophages and are both strongly connected to GSE2004, which contains samples from kidney, liver, and spleen, sites of long-lived macrophages. The first hematopoietic outlier, GSE2018 studies bronchoalveolar lavage cells from lung transplant receipts; the majority of these cells are macrophages. The dataset has strong links to solid tissue datasets including GSE2004, and the diverse dataset E-MEXP-66. The second hematopoietic outlier, GSE2665, is also strongly connected to GSE2004 and measures expression of lymphatic organs (sentinel lymph node) that contain sinusoidal macrophages and sinusoidal endothelial cells. The third outlier, E-MEXP-101, studies a colon carcinoma cell line and has connections to other cancer datasets in cluster 1.

Top dataset links overlap well with citation graph. We compared the model-driven network to the actual citation links (Fig. 2, right) to find out to what extent the citation practice in the research community matches the data-driven relationships. Of the top two hundred data-driven edges, 50% overlapped with direct or indirect citation links (see *Methods* and *SI Text*). Most of the direct citations appear within

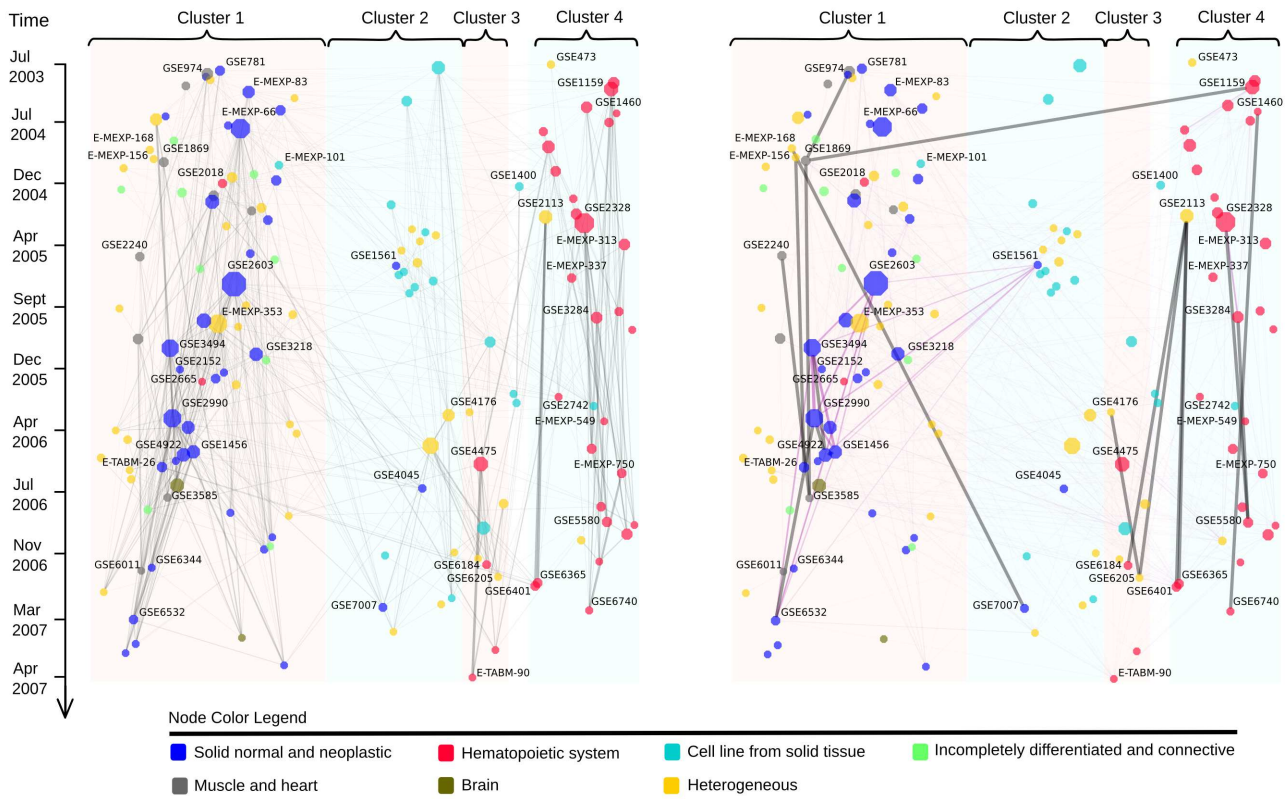

Fig. 2. Relevance network of datasets in the human gene expression atlas; data-driven links from the model (left) and citation links (right). Left: each dataset was used as a query to retrieve earlier datasets; a link from an earlier dataset to a later one means the earlier dataset is relevant as a partial model of activity in the later dataset. Link width is proportional to the normalized relevance weight (combination weight θ_j^q; only links with $\theta_j^q \geq 0.025$ are shown, and datasets without links have been discarded). Right: links are direct (gray) and indirect (purple) citations. Node size is proportional to the estimated influence, *i.e.*, the total outgoing weight. Colors: tissue types (six meta tissue types [12]). The node layout was computed from the data-driven network (details in *Methods*).

the four tissue clusters (Fig. 2, right). The two cross-cluster citations are not due to biological similarity of the datasets. The publication for GSE1869 cites the publication for GSE1159 regarding the method of differential expression detection. The GSE7007, a study on Ewing sarcoma samples, cites the study on human mesenchymal stem cells (E-MEXP-168) for stating that the overall gene expression profiles differ between those samples.

We additionally compared the densely connected sets of experiments between the two networks. In the citation graph the breast cancer datasets GSE2603, GSE3494, GSE2990, GSE4922, and GSE1456 form an interconnected clique in cluster 1, while the three leukocyte datasets GSE2328, GSE3284, and GSE5580 form an interconnected module in cluster 4. In the relevance network the corresponding edges for both cliques are among the strongest links for those datasets, and some of them are among the top 20 strongest edges in the network (see *SI Text* for the list of top 20 edges). There are also densely connected modules in the relevance network that are not strongly connected in the citation graph; when we systematically sought cliques associated to each of the top 20 edges, the most strong edges constitute a clique among E-MEXP-750, GSE6740 and GSE473, all three studying CD4+ T helper cells which are an essential part of the human immune system. Another interesting set is among three T-cell related datasets in cluster 3. Two of the datasets contain T lymphoblastic leukemia samples (E-MEXP-313 and E-MEXP-549), whereas E-MEXP-337 reports thymocyte profiles. Thymocytes are developing T lymphocytes that are matured in thymus, so this connection is biologically meaningful but not straightforward to find from dataset annotations. Other strongly connected cliques are discussed in the *SI Text*.

Analysis of network hubs discovers datasets deserving more citations. Datasets that have high weights in explaining other datasets have a large weighted outdegree in the data-driven relevance network, and are expected to be useful for many other studies. We checked whether the publications corresponding to these *central hubs* are highly cited in the research community. There is a low but statistically significant correlation between the weighted outdegree of datasets and their citation counts (Fig. 3; Spearman $\rho(169) = 0.2656$, $p < 0.001$). Both quantities were normalized to avoid bias due to different release times of the datasets (*Methods*). We further examined whether the prestige of the publication venue (measured by impact factor) and the senior author (h-index of the last author) biased the citation counts, which could explain the low correlation between the outdegree and the citation count, and the answer was affirmative (*Methods*).

We inspected more closely the datasets where the recommended or the actual citation counts were high (Fig. 3): (A) datasets having low citation counts but high outdegrees, (B) both high citation counts and high outdegrees and (C) high citation counts but low outdegrees. We manually checked the publication records of region A in Gene Expression Omnibus (GEO) [33] and ArrayExpress [17], to find out why the datasets had low citation counts despite their high outdegree (data-driven citation recommendations). Two of the eight datasets had an inconsistent publication record. The blue arrows in Fig. 3 point from their original position to the corrected position confirmed by GEO and ArrayExpress. Thus the data-driven network revealed the inconsistency, and the new positions, corresponding to higher citation counts, validate the model-based finding that these datasets are good explainers for other datasets. In region B, most of the papers have been published in high impact journals and have relatively high number of samples (average sample size of 154) compared to region

A (average sample size of 75). One of the eight datasets in the collection is the well known Connectivity Map experiment (GSE5258). Lastly the set C mostly contains unique targeted studies; there are five studies in the set, which are about leukocytes of injured patients, Polycomb group (PcG) proteins, senescence, Alzheimer's disease, and effect of cAMP agonist forskolin, a traditional Indian medicine. The studies have been published in high impact forums, and a possible reason of their low outdegree is their specific cellular responses, which are not very common in the atlas.

Discussion

Our main goal was to test the feasibility of the scenario where researchers let the data speak for themselves when relating new research to earlier studies. The conclusion is positive: even a relatively straightforward and scalable mixture modeling approach found both expected relationships such as tissue types, and relationships not easily found with keyword searches, including cells in different developmental stages or treatments resembling conditions in other cell types. While biologists could find such connections by bringing expert knowledge into keyword searches, the ultimate advantage of the data-driven approach is that it also yields connections beyond current knowledge, giving rise to new hypotheses and follow-up studies. For example, it seems surprising that the skeletal muscle dataset GSE6011 is linked also to kidney and brain datasets. Closer inspection yielded possible partial explanations. Some kidney areas are rich in blood vessels, lined by smooth muscle. Studies have shown common gene signatures between skeletal muscle and brain. Abnormal expression of the protein dystrophin leads to Duchenne muscular dystrophy, exhibited by a majority of samples in GSE6011; the brain is another major expression site for dystrophin [30]. Interestingly the top three potentially novel datasets, where only less than 50% of the expression pattern is modelled by earlier datasets (i.e. $\theta^q_{N_S+1} > 0.5$), are GSE2603 (a central breast cancer set), the Connectivity Map data (GSE5258) and the Burkitt's Lymphoma set (GSE4475; a cancer fundamentally distinct from other types of lymphoma). The first two are also recovered by the citation data (have relatively high citation counts and appear in region B in Fig. 3), unlike the third (which is part of region A in Fig. 3).

Our case study focused on global analysis of the relevance network obtained for a representative dataset collection, allowing for comparisons with the citation graph. The data-driven relationships corresponded to actual citations when available, but were richer and were able to spot out errors in citation links. Another intended use of the retrieval method is to support researchers in finding relevant data on a particular topic of interest. We performed a study to obtain insights into relationships among skeletal muscle datasets as well as between skeletal muscle and other datasets, and showed that the retrieval method lessens the need for laborious manual searches (*SI Text* and Fig. S4).

In this work we made simplifying assumptions: we only employed two model families, included biological knowledge only as pre-chosen gene sets, and assumed all new experiments to be mixtures of earlier ones, instead of sums of effects in them. We expect results to improve considerably with more advanced future alternatives, with the research challenge being to maintain scalability. Generalizability of the search across measurement batches, laboratories, and measurement platforms is a challenge. Our feasibility study showed that for carefully preprocessed datasets (of the microarray atlas [12]), data-driven retrieval is useful even across laboratories. Our method is generally applicable to any single platform, and takes into account the expert knowledge built into models of datasets for that platform; abstraction-based data representations, such as the gene set enrichment representation we used, have potential to facilitate cross-platform analysis. As data integration approaches develop further [34, 35], it may be possible to do searches even across different omics types; here, integration of meta data (pioneered in a specific semi-supervised framework [36]), several ontologies (MGED ontology, experimental factor ontology and ontology of biomedical investigations [37]) and text mining results [38, 39] are obviously useful first steps.

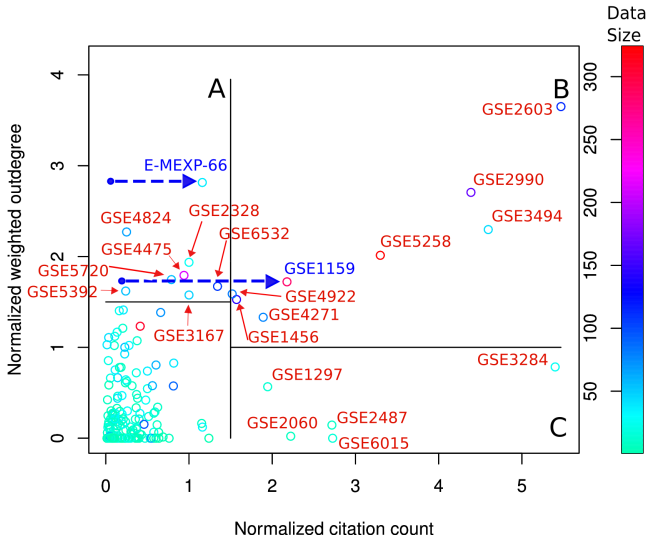

Fig. 3. Data-driven prediction of usefulness of datasets vs. their citation counts. Manual checks comparing sets for which the two scores differed revealed inconsistent database records for two datasets; the blue arrows point to their corrected locations, which are more in line with the data-driven model. Regions A, B, and C: see text.

Materials and Methods

Gene expression data. We used the human gene expression atlas [12] available at ArrayExpress under accession number E-MTAB-62. The data were preprocessed by gene set enrichment analysis (GSEA) using the canonical pathway collection (C2-CP) from the Molecular Signatures Database [19]. Each sample was represented by its top enriched gene sets [20] (*SI Text*).

Node layout and normalized relevance weight. The weight matrix contains a weight vector for each query dataset, encoding the amount of variation in that query explained by each earlier dataset. As query datasets from early years have only few even earlier sets available, there is a bias towards the edges being stronger for the datasets from early years. To remove the bias we normalized, for the visualizations, the edge strengths of each query data set by the number of earlier datasets. To visualize the relationship network over time in Fig. 2, we needed a layout algorithm that positions the datasets on the horizontal axis highlighting structure and avoiding tangling. We used a *cluster-emphasizing* Sammon's mapping [40]; Sammon's mapping is a nonlinear projection method or Multidimensional Scaling algorithm which aims at preserving the interpoint distances (here $1 - \theta^q_j$). By clustering the network (with unsupervised Markov clustering [41]) and increasing between-cluster distances by adding a constant ($c = 1$) to them, the mapping was made to emphasize clusters and hence untangle the layout.

Citation graph. Direct citations between dataset-linked publications were extracted from the Web of Science (26 Jul 2012) and PubMed (17 Oct 2012). We additionally considered two types of indirect edges. Firstly, we introduced links between datasets whose publications share common references. This covers for instance related datasets whose publications appeared close in time, making direct citation unlikely. A natural measure of edge strength is given by the number of shared references. Secondly, we connect datasets whose articles are cited together, because co-citation is a sign that the community perceives the articles as related. Here, the edge strength was taken to be the number of articles co-citing the two dataset publications; these edges dominate the indirect links in the citation graph. For this analysis we used citation data, available for 171 datasets and provided by Thomson Reuters as of 13 September 2012.

Normalization of citation counts and weighted outdegrees. As early datasets have many more papers which can cite them, and many more later datasets which they can help model, both the citation counts and estimated

weighted outdegrees are expected to be upwards biased for them. For Fig. 3 we normalized the quantities; for each dataset we normalized the outdegree by the number of newer datasets, and the citation count by the time difference between publishing the data and the newest dataset in the atlas. To make sure the normalization did not introduce side effects we additionally checked that the same conclusions were reached without the citation count normalization (Fig. S1; plotted as stratified subfigures for each 1-year time window). The citation counts were extracted from PubMed on 16 May 2012.

Citation counts are strongly influenced by external esteem of the publication forum and the senior author. We stratified the data sets according to the numbers of data-driven citation recommendations, and studied whether the impact factor of the forum or the h-index of the last author were predictive of the actual citation count in each stratum. The strata were the top and bottom quartiles, and for each we compared the top and bottom quartiles of the actual citation counts (resulting in comparing the four corners of Fig. 3). For low outdegree (low recommended citation count), the h-index was lower for less cited datasets ($t_{11} = 2.78, p = 0.0086$; mean value 24.20 vs 54.62), and also the impact factor was lower ($t_7 = 2.6, p = 0.016$; mean value 4.38 vs 21.13). Similarly, for high recommended citation count the impact factor for the little-cited datasets was lower ($t_{19} = 3.99, p = 4.0^{-4}$; mean value 6.45 vs 21.91), while the difference in h-index was not significant. All t statistics and p-values were computed by one-sided independent sample Welch's t-tests. The h-indices and impact factors were collected from Thomson Reuters Web of Knowledge and Journal Citation Reports 2011 respectively on 23rd July 2012.

ACKNOWLEDGMENTS. We thank Matti Nelimarkka and Tuukka Ruotsalo for helping with citation data. Certain data included herein are derived from the following indices: Science Citation Index Expanded, Social Science Citation Index and Arts & Humanities Citation Index, prepared by Thomson Reuters®, Philadelphia, Pennsylvania, USA, © Copyright Thomson Reuters ®, 2011. This work was financially supported by the Academy of Finland (Finnish Centre of Excellence in Computational Inference Research COIN, grant no 251170).

1. Greene C-S, Troyanskaya O-G (2011) PILGRM: An interactive data-driven discovery platform for expert biologists. *Nucleic Acids Res* 39:W368–74.
2. Tanay A, Steinfeld I, Kupiec M, Shamir R (2005) Integrative analysis of genome-wide experiments in the context of a large high-throughput data compendium. *Mol Syst Biol* 1:e1–10.
3. Caldas J, et al. (2012) Data-driven information retrieval in heterogeneous collections of transcriptomics data links *SIM2s* to malignant pleural mesothelioma. *Bioinformatics* 28:i246–i253.
4. Adler P, et al. (2009) Mining for coexpression across hundreds of datasets using novel rank aggregation and visualization methods. *Genome Biol* 10:R139.
5. Schmid P-R, Palmer N-P, Kohane I-S, Berger B (2012) Making sense out of massive data by going beyond differential expression. *Proc Natl Acad Sci U S A* 109:5594–5599.
6. Gerber G-K, Dowell R-D, Jaakkola T-S, Gifford D-K (2007) Automated discovery of functional generality of human gene expression programs. *PLoS Comput. Biol* 3:e148.
7. Tseng G-C, Ghosh D, Feingold E (2012) Comprehensive literature review and statistical considerations for microarray meta-analysis. *Nucleic Acids Res* 40:3785–3799.
8. Rung J, Brazma A (2012) Reuse of public genome-wide gene expression data. *Nat Rev Genet* 14:89–99.
9. Baxter J (1997) A Bayesian/information theoretic model of learning to learn via multiple task sampling. *Machine Learning* 28:7–39.
10. Caruana R (1997) Multitask learning. *Machine Learning* 28:41–75.
11. Punta M, et al. (2012) The Pfam protein families database. *Nucleic Acids Res* 40:D290–D301.
12. Lukk M, et al. (2010) A global map of human gene expression. *Nat Biotechnol* 28:322–324.
13. Russ J, Futschik M-E (2010) Comparison and consolidation of microarray data sets of human tissue expression. *BMC Genomics* 11:305.
14. Suthram S, et al. (2010) Network-based elucidation of human disease similarities reveals common functional modules enriched for pluripotent drug targets. *PLoS Comput Biol* 6:e1000662.
15. Huttenhower C, Troyanskaya O-G (2008) Assessing the functional structure of genomic data. *Bioinformatics* 24:i330–8.
16. Meinicke P, Asshauer K-P, Lingner T (2011) Mixture models for analysis of the taxonomic composition of metagenomes. *Bioinformatics* 27:1618–24.
17. Parkinson H, et al. (2009) ArrayExpress update–from an archive of functional genomics experiments to the atlas of gene expression. *Nucleic Acids Res* 37:D868–72.
18. Gionis A, Indyk P, Motwani R (1999) Similarity search in high dimensions via hashing. *Proceedings of the Twenty-Fifth International Conference on Very Large Databases* eds Atkinson M, Orlowwska M, Valduriez P, Zdonik S, Brodie M (Edinburgh, Scotland) pp 518–529.
19. Subramanian A, et al. (2005) Gene set enrichment analysis: A knowledge-based approach for interpreting genome-wide expression profiles. *Proc Natl Acad Sci U S A* 102:15545–15550.
20. Caldas J, Gehlenborg N, Faisal A, Brazma A, Kaski S (2009) Probabilistic retrieval and visualization of biologically relevant microarray experiments. *Bioinformatics* 25:i145–i153.
21. Engreitz J-M, et al. (2010) Content-based microarray search using differential expression profiles. *BMC Bioinformatics* 11:603.
22. Pritchard J-K, Stephens M, Donnelly P (2000) Inference of population structure using multilocus genotype data. *Genetics* 155:945–959.
23. Blei D-M, Ng A-Y, Jordan M-I, Lafferty J (2003) Latent Dirichlet allocation. *J Mach Learn Res* 3:993–1022.
24. Nigam K, McCallum A, Thrun S, Mitchell T (2000) Text classification from labeled and unlabeled documents using EM. *Machine Learning* 39:103–134.
25. Zhu Y, Davis S, Stephens R, Meltzer P-S, Chen Y (2008) GEOmetadb: powerful alternative search engine for the Gene Expression Omnibus. *Bioinformatics* 24:2798–2800.
26. Mourkioti F, et al. (2013) Role of telomere dysfunction in cardiac failure in Duchenne muscular dystrophy. *Nature Cell Bio* 15:895–904.
27. Kitazawa M, Trinh D-N, LaFerla F-M (2008) Inflammation induces tau pathology in inclusion body myositis model via glycogen synthase kinase-3 beta. *Ann Neurol* 64:15–24.
28. Martinsson L, et al. (2013) Long-term lithium treatment in bipolar disorder is associated with longer leukocyte telomeres. *Transl Psychiatry* 3:e261.
29. Kimura K, et al. (1995) Diversity and variability of smooth muscle phenotypes of renal arterioles as revealed by myosin isoform expression. *Kidney Int* 48:372–382.
30. Culligan K, Glover L, Dowling P, Ohlendieck K (2001) Brain dystrophin-glycoprotein complex: Persistent expression of beta-dystroglycan, impaired oligomerization of Dp71 and up-regulation of utrophins in animal models of muscular dystrophy. *BMC Cell Biol* 2:2.
31. Kirchner T, et al. (1988) Pathogenesis of myasthenia gravis. acetylcholine receptor-related antigenic determinants in tumor-free thymuses and thymic epithelial tumors. *Am J Pathol* 130:268–280.
32. Holliday D-L, Speirs V (2011) Choosing the right cell line for breast cancer research. *Breast Cancer Res* 13:215.
33. Barrett T, et al. (2011) NCBI GEO: archive for functional genomics data sets-10 years on. *Nucleic Acids Res* 39:D1005–D1010.
34. Tripathi A, Klami A, Orešič M, Kaski S (2011) Matching samples of multiple views. *Data Min Knowl Discov* 23:300–321.
35. Virtanen S, Klami A, Khan S-A, Kaski S (2012) Bayesian group factor analysis. *JMLR Workshop Conf Proc* 22:1269–1277.
36. Wise A, Oltvai Z, Bar-Joseph Z (2012) Matching experiments across species using expression values and textual information. *Bioinformatics* 28:i258–i264.
37. Zheng J, et al. (2011) Annotcompute: annotation-based exploration and meta-analysis of genomics experiments. *Database (Oxford)*.
38. Jensen L-J, Saric J, Bork P (2006) Literature mining for the biologist: from information retrieval to biological discovery. *Nat Rev Genet* 7:119–129.
39. Rzhetsky A, Seringhaus M, Gerstein M (2008) Seeking a new biology through text mining. *Cell* 134:9–13.
40. Sammon J-W (1969) A nonlinear mapping for data structure analysis. *IEEE Trans Comput* 18:401–409.
41. van Dongen S (2000) Graph Clustering by Flow Simulation. *PhD thesis, University of Utrecht*.

arXiv:1404.0329v1 [q-bio.QM] 1 Apr 2014

Supporting Information Appendix: Toward computational cumulative biology by combining models of biological datasets

Ali Faisal [*], **Jaakko Peltonen** [*], **Elisabeth Georgii** [*], **Johan Rung** [†] **and Samuel Kaski** [*] [‡]

[*]Helsinki Institute for Information Technology HIIT, Department of Information and Computer Science, PO Box 15400, Aalto University, 00076 Aalto, Finland,[†]European Molecular Biology Laboratory, European Bioinformatics Institute (EMBL-EBI), Wellcome Trust Genome Campus Hinxton, CB10 1SD, UK, and [‡]Helsinki Institute for Information Technology HIIT, Department of Computer Science, PO Box 68, 00014 University of Helsinki, Finland

Methods

Gene Set Enrichment Analysis. We used GSEA [1] to bring in biological knowledge in the form of pre-defined gene sets. GSEA starts by sorting genes with respect to their normalized expression levels. GSEA essentially consists of computing a running sum on the sorted list for each gene set; this running sum (enrichment score) increases when a gene belongs to the gene set and decreases otherwise; the final statistic is the maximum of this running sum. The procedure essentially amounts to computing a weighted Kolmogorov-Smirnov (KS) statistic. For each sample the KS statistic is normalized by dividing it by the mean of random KS statistics computed on randomly generated gene sets whose size is matched with the actual gene set; the 50 top-scoring gene sets are selected according to this normalized score. This simple thresholding ignores significance values but has been successfully used in earlier meta-analysis studies [2, 3, 4]; we earlier investigated the alternative of selecting gene sets based on a standard q-value cut-off (with $q < 0.05$), but it produced an excessively sparse encoding where more than 80% samples had no active gene sets [4]. The activity of each gene set is finally expressed as its core set or *leading edge subset*, consisting of genes found before the running KS score reached its maximum. We quantify the activity or a set simply by the size of the leading edge subset. It can be used analogously to the so called *bag-of-words* representations in text analysis, and we use it for the subsequent modeling of each dataset with the base models.

Base models. *Latent Dirichlet allocation* (LDA) [5, 6, 7] and *mixture of unigrams* [8] are probabilistic unsupervised models that give insight to datasets by describing them in terms of latent components. Each data sample, in our case quantified as gene set activities, is represented by a probability distribution over components (sometimes also called topics). The components are shared by all samples, but with different degrees of activation for each, and each component produces a characteristic distribution over the gene sets. In LDA each sample may be produced by multiple hidden components while the mixture of unigrams is a simplified version where each sample is assumed to come from a single component. The computational problem is to estimate the latent component structure (for each sample, the distribution over components and for each component, its distribution over gene sets) that has most likely generated the observed set of samples.

We use standard inference solutions for the models: collapsed Gibbs sampler for the LDA [7] and Expectation Maximization for the mixture of unigrams [8]. In LDA, the hyperparameters, the prior probability of each component, were optimized with Minka's stable fixed point iteration scheme [10], interleaved with the collapsed Gibbs sampling of the other parameters. The number of Gibbs iterations was 2500, found to yield good performance also earlier [6, 7, 9]. For the mixture of unigrams model, the maximum a posterior solution is estimated with the Expectation Maximization (EM) algorithm, using Laplace smoothing for the prior probabilities of the mixture components [8].

Model selection for base models. For each dataset we estimated the two base models and selected the one that best models the dataset. We split the dataset into two parts where 90% of the dataset samples were used for training the two models and the remaining 10% samples to compute the test-set predictive likelihood, a measure of how well the model fits the data. We repeated this procedure in a 10-fold cross-validation setup for each dataset and each of the two models. The model that performed better on average across the 10 folds was chosen to represent the dataset. Most datasets (74 out of 112) preferred the more expressive LDA model.

The number of components was selected with cross-validation for both models. We again used 10-fold cross-validation, separately for each dataset, to choose the number of components that lead to best predictive likelihoods on test samples from the same dataset. We observed that the optimal number of total topics could roughly be summarized by $\lceil N_D/3 \rceil$ for LDA and $\lfloor \sqrt{N_D} \rfloor$ for mixture of unigrams, where N_D is the total number of samples in a dataset. All predictive likelihoods (both for model selection and for retrieval) were computed using an empirical likelihood method (see [11]). For very small datasets ($N_D < 10$) within-set cross-validation is very noisy and we chose LDA which had been selected for the majority of datasets with $10 \leq N_D \leq 15$.

Strict Concavity of the objective function. For each query dataset q, there are multiple samples $i = 1, \ldots, N_q$. For each sample, each background dataset gives a probability. Let $\vec{x}_i = [p(x_i^q|M^{s_1}), p(x_i^q|M^{s_2}), \ldots, p(x_i^q|M^{s_{N_S}}), p(x_i^q|\Psi)]^\top$ be the column vector of probabilities for sample i from all background datasets and from the novelty model, where $^\top$ denotes vector transpose. Then the probability given by a mixture of background datasets is $\vec{\theta}^\top \vec{x}_i$ where $\vec{\theta}$ is the column vector of mixture weights (mixing proportions) over the background datasets and the novelty model. Since the mixing weights are mixture probabilities, they must lie in the canonical simplex denoted as

$$\Delta = \left\{ \theta | \theta_j \geq 0, \sum_{j=1}^{N_S+1} \theta_j = 1 \right\} . \qquad \textbf{[1]}$$

Our optimization takes the form

$$\max_{\vec{\theta} \in \Delta} \log \left(\exp(-\lambda ||\vec{\theta}||^2) \prod_i \vec{\theta}^\top \vec{x}_i \right) = \max_{\vec{\theta} \in \Delta} \sum_i \log(\vec{\theta}^\top \vec{x}_i) - \lambda ||\vec{\theta}||^2 = \max_{\vec{\theta} \in \Delta} f(\vec{\theta}) \qquad \textbf{[2]}$$

where the objective function is

$$f(\vec{\theta}) = \sum_i \log(\vec{\theta}^\top \vec{x}_i) - \lambda ||\vec{\theta}||^2 .$$

The objective $f(\vec{\theta})$ is a strictly concave function with respect to the multivariate parameter $\vec{\theta}$. This can be shown since the second

derivative of the function is negative. In detail, the function's gradient is

$$\nabla f(\vec{\theta}) = \sum_i \frac{\vec{x}_i}{\vec{\theta}^\top \vec{x}_i} - 2\lambda\vec{\theta} \quad \textbf{[3]}$$

and the matrix of second-order partial derivatives (Hessian matrix) is

$$\nabla^2 f(\vec{\theta}) = -\sum_i \frac{\vec{x}_i \vec{x}_i^\top}{(\vec{\theta}^\top \vec{x}_i)^2} - 2\lambda I \quad \textbf{[4]}$$

where I denotes the identity matrix. Since all elements of $\vec{x}_i$ are non-negative, it is easy to see the Hessian matrix is negative definite for all $\vec{\theta}$ where $\theta_i \geq 0$ as long as $\lambda > 0$. Therefore the objective function is strictly concave. A local maximum of a strictly concave function on a convex feasible region is the unique global maximum [12]; therefore maximizing the objective function to a local maximum by any algorithm yields the unique global maximum.

Maximizing the Objective Function. Maximization of concave functions on the unit simplex Δ can be done by the *Frank-Wolfe algorithm*. The algorithm performs the following steps.

Step 1: initialization. The algorithm initializes a solution as the vertex of the simplex having the largest objective value. A vertex of the simplex has $\theta_j = 1$ for some j and all other elements of $\vec{\theta}$ are zero. At vertex j, the objective function simplifies to

$$\sum_i \log(x_{ij}) - \lambda$$

where $x_{ij} = p(sample_i|dataset_j)$, which takes $O(N_q)$ time to evaluate per vertex. Thus creating the initialization takes linear $O(N_q N_S)$ time even with a simple brute-force evaluation of all vertices.

Step 2: Iteration. In each iteration, the algorithm improves the solution by two steps: (1) Find the maximal element j of the gradient. (2) Find the point along the line $\vec{\theta} + \alpha(\vec{e}(j) - \vec{\theta})$, $\alpha \in [0,1]$, which maximizes the objective function. Here $\vec{e}(j)$ means the vector where the element j is one and the others are zero. Computation of the gradient and finding the maximal element takes $O(N_q N_S)$ time.

Proof of Convergence and Scalability

Convergence Analysis. Since each iteration takes linear time with respect to the number of query samples and background datasets, the only remaining issue is the number of iterations required for good enough convergence. We now analyze the convergence properties of this iteration in two cases.

Case 1: optimal α. At first, we consider the case where the best value of α can be found along the line to a sufficient accuracy in a fixed amount of time, for example by restricting evaluations along the line to a fixed number.

Define a proportional regret function $h(\vec{\theta}) = (f(\vec{\theta}^*) - f(\vec{\theta}))/4C_f$ where $\vec{\theta}^*$ is the optimal parameter value maximizing the objective function, and $C_f \geq 0$ is a *measure of curvature* of f. In detail, C_f is defined as the largest quantity such that for all $\vec{\theta}_A \in \Delta$, $\vec{\theta}_B \in \Delta$, $\vec{\theta}_C$ where $\vec{\theta}_C = \vec{\theta}_A + \alpha(\vec{\theta}_B - \vec{\theta}_A)$ for some α, we have

$$f(\vec{\theta}_C) \geq f(\vec{\theta}_A) + (\vec{\theta}_C - \vec{\theta}_A)^\top \nabla f(\vec{\theta}_A) - \alpha^2 C_f \, .$$

With this notation, it can be shown ([13], Theorem 2.2) that at iteration k of the Frank-Wolfe algorithm the current solution $\vec{\theta}_k$ has regret

$$h(\vec{\theta}_k) \leq 1/(k+3)$$

and thus

$$f(\vec{\theta}^*) - f(\vec{\theta}_k) \leq 4C_f/(k+3) \, .$$

Thus, to achieve a desired regret ε, at most $4C_f/\varepsilon + 3$ iterations are needed independently of the number of background datasets; the amount of iterations needed depends only on the curvature.

Case 2: fixed α. It can even be shown that using a fixed value $\alpha_k = 2/(k+3)$ at each iteration k suffices to yield bounds for performance: with this choice of α_k the regret bound at iteration $k+1$ becomes ([13], Section 7)

$$h(\vec{\theta}_{k+1}) \leq 1/(k+4)$$

and thus

$$f(\vec{\theta}^*) - f(\vec{\theta}_{k+1}) \leq 4C_f/(k+4)$$

which again shows the number of iterations to achieve a desired regret does not depend on the number of background datasets, only on the curvature. It is enough to show the curvature is finite and does not depend on the number of background datasets; we now show this.

Analysis of the Curvature: As seen above, the smaller the curvature C_f, the better the bounds for the regret $f(\vec{\theta}^*) - f(\vec{\theta}_k)$ are. In our case the function f is twice differentiable and it can be shown ([13], Section 4.1) that

$$C_f \leq \sup_{\vec{\theta}_A \in \Delta, \vec{\theta}_B \in \Delta, \alpha \in [0,1]} -\frac{1}{2}(\vec{\theta}_B - \vec{\theta}_A)^\top \nabla^2 f(\vec{\theta}_\alpha)(\vec{\theta}_B - \vec{\theta}_A) \, .$$

where $\vec{\theta}_\alpha = \vec{\theta}_A + \alpha(\vec{\theta}_B - \vec{\theta}_A)$. For our cost function this becomes

$$\begin{aligned}
C_f &\leq \sup_{\vec{\theta}_A \in \Delta, \vec{\theta}_B \in \Delta, \alpha \in [0,1]} \frac{1}{2}(\vec{\theta}_B - \vec{\theta}_A)^\top \\
&\qquad \left(\sum_i \frac{\vec{x}_i \vec{x}_i^\top}{(\vec{\theta}_\alpha^\top \vec{x}_i)^2} + 2\lambda I \right) (\vec{\theta}_B - \vec{\theta}_A) \\
&= \sup_{\vec{\theta}_A \in \Delta, \vec{\theta}_B \in \Delta, \alpha \in [0,1]} \frac{1}{2} \\
&\qquad \left(\sum_i \frac{((\vec{\theta}_B - \vec{\theta}_A)^\top \vec{x}_i)^2}{(\vec{\theta}_\alpha^\top \vec{x}_i)^2} + 2\lambda ||\vec{\theta}_B - \vec{\theta}_A||^2 \right) \\
&\leq \sup_{\vec{\theta}_A \in \Delta, \vec{\theta}_B \in \Delta, \alpha \in [0,1], 1 \leq i \leq N_q} \frac{1}{2} \left(N_q \frac{((\vec{\theta}_B - \vec{\theta}_A)^\top \vec{x}_i)^2}{(\vec{\theta}_\alpha^\top \vec{x}_i)^2} + 4\lambda \right) \\
&\leq \sup_{\vec{\theta}_A \in \Delta, \vec{\theta}_B \in \Delta, \alpha \in [0,1], 1 \leq i \leq N_q} \frac{1}{2} \left(N_q \frac{(\vec{\theta}_B \vec{x}_i)^2 + (\vec{\theta}_A^\top \vec{x}_i)^2}{(\vec{\theta}_\alpha^\top \vec{x}_i)^2} + 4\lambda \right) \\
&\leq \sup_{\vec{\theta}_A \in \Delta, \vec{\theta}_B \in \Delta, \alpha \in [0,1], 1 \leq i \leq N_q} \frac{1}{2} \left(N_q \frac{2(\max_j x_{ij})^2}{(\min_j x_{ij})^2} + 4\lambda \right) \quad \textbf{[5]}
\end{aligned}$$

where for brevity we denoted $x_{ij} = p(x_i^q|M^{s_j})$ for $j = 1, \ldots, N_S$ and $x_{i,N_S+1} = p(x_i^q|\Psi)$. Notice that the right-hand side only depends on the maximal and minimal values that background datasets give to query samples, not on the number of such background datasets. Thus, as long as we ensure the maximal value is below some finite constant and the minimal value is above some small nonzero constant, the curvature C_f is finite and the convergence bounds of the Frank-Wolfe algorithm therefore do not depend on the number of background datasets. This condition is simple to ensure by suitable regularization of the models of background datasets.

Therefore under the simple condition that probabilities given to query samples are upper bounded and lower bounded above zero, the algorithm converges (towards the unique global maximum) to a desired tolerance of the regret in a finite number of iterations I, which can depend on the number of query samples but is independent of the size of number of background datasets.

Computational complexity. Our model needs to perform two main computation tasks for each new dataset: optimization of the objective function and computation of the predictive likelihoods. The optimization step needs to evaluate the function value, the gradient and maximal element of the gradient; the computation of the function value has complexity $O(N_q * N_S * O(\text{computing } p(x_i{}^q|M^{s_j}))$ while, as discussed in the previous section, computing the gradient and finding the maximal element takes linear $O(N_q * N_S)$ time in each iteration if a fixed step size is used (or if the number of line search evaluations is restricted below some fixed maximum). Thus the complete algorithm takes $O(I * N_q * N_S * O(\text{computing } p(x_i{}^q|M^{s_j})) + O(I * N_q * N_S)$. Clearly the dominating factor is the computation of the function value: $O(I * N_q * N_S * O(\text{computing } p(x_i{}^q|M^{s_j}))$.

The predictive likelihood for the query sample is computed as its average probability from $V = 1000$ multinomial distributions, estimated from randomly generated samples coming from the generative process of the earlier-trained model M^{s_j}; this is the standard *empirical likelihood scheme* discussed in [11]. The complexity of computing the predictive likelihood for the query dataset with G features (in our case gene-sets), given a model with T latent components [1] is $O(G * V * T)$. The total computational complexity is then simply the complexity for predictive likelihoods times the optimization scheme: $O(I * N_q * N_S * G * V * T)$. Since the number of iterations is independent of the number of background datasets, N_S, the complexity is linear with respect to it, $O(N_S)$, and therefore the model is reasonably tolerant to the fast growth of public repositories.

In our implementation a single query dataset took about 31 iterations and 0.15 seconds on average on an Intel® Core (TM) i7 CPU @ 2.93GHz, to find the optimal weight vector.

Results

Normalization of citation counts. Older datasets tend to have higher citation counts and outdegrees; in Fig. 3 we removed this bias by a normalization technique, and identified interesting datasets having very low citation counts and very high outdegrees or vice versa. To verify the analysis results (identified datasets) are not artifacts of the normalization, we reanalyzed the original data without applying citation count normalization. Fig. S1 shows the result as stratified subfigures plotted for each year separately. The datasets identified using the original citation counts (top-left and bottom-right corners in each subfigure of Fig. S1) are an exact match with the datasets identified after normalization.

Retrieval performance after discounting for the laboratory effect. In microarray experiments laboratory effects are known to be strong [14]. The 206 datasets studied were generated from 163 laboratories. The top laboratory was responsible for 7 datasets, whereas 142 laboraties only contributed a single set. To test how much the laboratory effects have affected our results, we discarded all retrieved results from the same laboratory as the query set. The original precision-recall curve and the corrected curve are in Figure S2; the mean average precision dropped to 0.44 from 0.42. The small change in performance shows that our result is mainly due to other effects captured by the model than the laboratory effect.

Quantitative comparison of data-driven results against the citation patterns. Of the 23 direct citation links, eleven were also found by our model as having a non-zero edge weight. Six links could not have been found; five of them are citations by papers having very small datasets (< 10), which we had considered to be too small to act as queries while one citation link is between datasets released on the same date. Of the remaining six links not observed in our model, two are the cross-cluster citations that are not due to biological similarity of the datasets as discussed in the paper; two cell line datasets about multiple myeloma and large cell lymphoma (GSE6205 and GSE6184 respectively) cite a leukemia study (GSE2113) where plasma cells are profiled; one dataset measuring HIV infection from T cells (GSE6740) cites a very small dataset about thymocyte, the limited size of the set reducing its corresponding model's relative capability to explain other sets compared to large datasets in the collection. Finally, E-TABM-26, a prostate cancer study, cites E-MEXP-156 (a study about tumorigenic and nontumorigenic Human Embryonic Kidney Cells) in the context of cell apoptosis of cancer cells in general.

Vice versa, we evaluated the top-weight data-driven edges against direct and indirect citation links and found a favorable, non-random overlap. For this we use a modern standard metric in information retrieval called precision *@k* which measures precision at each position *k* in the ranked list of top retrieval results from a search engine. In our case the results are the ranked list of inferred edges in descending order of their edge strengths. The citation patterns were used as the gold standard for existence of an edge between two datasets using their corresponding publication information. Fig. S3 shows that the precision *@k* is reasonably high; for the top two hundred data-driven edges (*i.e.*, $k = 200$) the value is 0.5.

Densely connected set of experiments in the relevance network. For each of the top 20 strongest edges in the relevance network (listed in Table S1) we searched for associated cliques, *i.e.*, connections within neighbors of the edge, where the clique size is at least three. We found seven cliques; four of them (breast cancer, leukocytes, human immune T helper cells and developmental stages of thymocytes) are described in the main text. The remaining three are an adenocarcinoma, a brain tissue and a skeletal muscle clique. The first clique is among GSE4824, GSE5258 and GSE6914; all three datasets are heterogeneous collections of different cancerous tissues where majority of the samples are cell line profiles from either lung or breast adenocarcinoma. The second clique is among GSE1297, GSE5392 and E-MEXP-114 that profile normal and diseased brain tissue. The last clique is among all skeletal muscle experiments of the collection where the strongest edge is between GSE3307 and GSE6011 that both measure Duchenne muscular dystrophy sampled from quadriceps muscle tissues, the former also containing samples form other skeletal muscle diseases.

Skeletal muscle retrieval case study. We lastly present a case study to illustrate how the model retrieval can support a researcher in finding relevant data on a specific topic, lessening the need for laborious manual searches. The topic of interest was gene expression in skeletal muscle and our database consisted of the human gene expression atlas [15], which includes eight skeletal muscle datasets among a total of 206 datasets, plus additional 16 skeletal muscle datasets extracted manually from ArrayExpress (Table S2). First, we tested how well the data-driven method retrieved other skeletal muscle datasets when querying with any single skeletal muscle dataset. The retrieval performance across all 16 query datasets was close to optimal, whereas keyword searches were not as good in the same task (Fig. S4). The reason for that was the lack of a consistent annotation vocabulary in the dataset descriptions; in particular, different levels of specificity had been used.

Next we looked in more detail into the retrieval results of the individual queries. For all of them, the retrieval result was sparse, *i.e.*, less than 10% of the datasets were found to be relevant to the query (by a non-zero weight). We further investigated the ranking of results provided by the data-driven retrieval model. For 12 queries all retrieved results were other skeletal muscle datasets; results for the remaining queries contained at least one false positive, and they are summarized in Table S3. The top retrieved non-skeletal muscle dataset is the brain tissue dataset GSE5392, followed by the kidney dataset GSE781. Kidney does not have a direct biological connection

[1] T is upper bounded by the maximum number of samples in a background dataset.

to skeletal muscle. It is known that some areas of kidney are rich in blood vessels, which are lined by smooth muscle. Skeletal muscle, smooth muscle and cardiac muscle are the three main muscle types in the human body. Interestingly, one skeletal muscle datasets, E-MEXP-216, was never retrieved by any skeletal muscle query, which suggests that it is an outlier in some respect. The dataset contains a combination of human and macacque liver and skeletal muscle samples. It contains only four human skeletal muscle samples.

Finally, the internal ranking of skeletal muscle datasets in response to a particular query dataset seems to be guided to a large extent by health conditions:

- **E-GEOD-9397** contains samples annotated with disease status FSHD (Facioscapulohumeral muscular dystrophy). The top retrieved dataset for that query is E-GEOD-10760, the only other dataset in the collection annotated with FSHD.
- The top retrieved result for **E-GEOD-12648** is **E-GEOD-11686** and vice versa. Both of them are neuromuscular disorders, hereditary inclusion body myopathy (HIBM) and cerebral palsy, and they are the only datasets with these specific diseases in our data collection.
- **E-GEOD-1786** partially contains samples from COPD (Chronic Obstructive Pulmonary Disease) subjects; the disease has a muscle wasting effect. The top retrieved result is E-GEOD-10760; it contains samples from the same muscle type, but another disease (FSHD) which also leads to muscle wasting.
- **E-GEOD-1295** contains samples of the trained and untrained muscle of non-young overweight people with prediabetic metabolic syndrome. Half of the samples in E-GEOD-1786 are also from trained muscles (of COPD patients and controls, old overweight men). E-GEOD-1786 appears at rank 6, and it is the only background dataset known to contain trained muscle samples. Among the top 5 datasets, 3 are annotated to contain disease samples related to weakening of muscles, another one is known to contain old overweight samples (**E-GEOD-8441**).

In summary, disease conditions and health states of tissue seem to determine the ranking within datasets of the same tissue (sub)type.

1. Subramanian A, et al. (2005) Gene set enrichment analysis: A knowledge-based approach for interpreting genome-wide expression profiles. *Proc Natl Acad Sci U S A* 102:15545–15550.
2. Segal E, et al. (2004) A module map showing conditional activity of expression modules in cancer. *Nat Genet* 36:1090–1098.
3. Caldas J, Gehlenborg N, Faisal A, Brazma A, Kaski S (2009) Probabilistic retrieval and visualization of biologically relevant microarray experiments. *Bioinformatics* 25:i145–i153.
4. Caldas J, et al. (2012) Data-driven information retrieval in heterogeneous collections of transcriptomics data links *SIM2s* to malignant pleural mesothelioma. *Bioinformatics* 28:i246–i253.
5. Pritchard J-K, Stephens M, Donnelly P (2000) Inference of population structure using multilocus genotype data. *Genetics* 155:945–959.
6. Blei D-M, Ng A-Y, Jordan M-I, Lafferty J (2003) Latent Dirichlet allocation. *J Mach Learn Res* 3:993–1022.
7. Griffiths T-L, Steyvers M (2004) Finding scientific topics. *Proc Natl Acad Sci U S A* 101:5228–5235.
8. Nigam K, Mccallum A-K, Thrun S, Mitchell T (2000) Text classification from labeled and unlabeled documents using EM. *Machine Learning* 39:103–134.
9. Teh Y-W, Jordan M-I, Beal M-J, Blei D-M (2006) Hierarchical Dirichlet Processes. *J Am Stat Assoc* 101:1566–1581.
10. Minka T-P (2000) Estimating a Dirichlet distribution. *Technical report* available at http://research.microsoft.com/en-us/um/people/minka/papers/dirichlet/.
11. Li W, McCallum A (2006) Pachinko allocation: DAG-structured mixture models of topic correlations. *Proceedings of the Twenty-Third international conference on Machine learning* (ACM New York, USA) pp 577–584.
12. Bradley S, Hax A, Magnanti T (1977) *Applied Mathematical Programming*. Addison-Wesley.
13. Clarkson, K L (2010) Coresets, sparse greedy approximation, and the Frank-Wolfe algorithm. *ACM Transactions on Algorithms* 6(4): Article 63.
14. Zilliox M-J, Irizarry R-A (2007) A gene expression bar code for microarray data. *Nat Methods* 4:911–913.
15. Lukk M, et al. (2010) A global map of human gene expression. *Nat Biotechnol* 28:322–324.
16. Shannon P, et al. (2003) Cytoscape: a software environment for integrated models of biomolecular interaction networks. *Genome Res* 13:2498-2504.

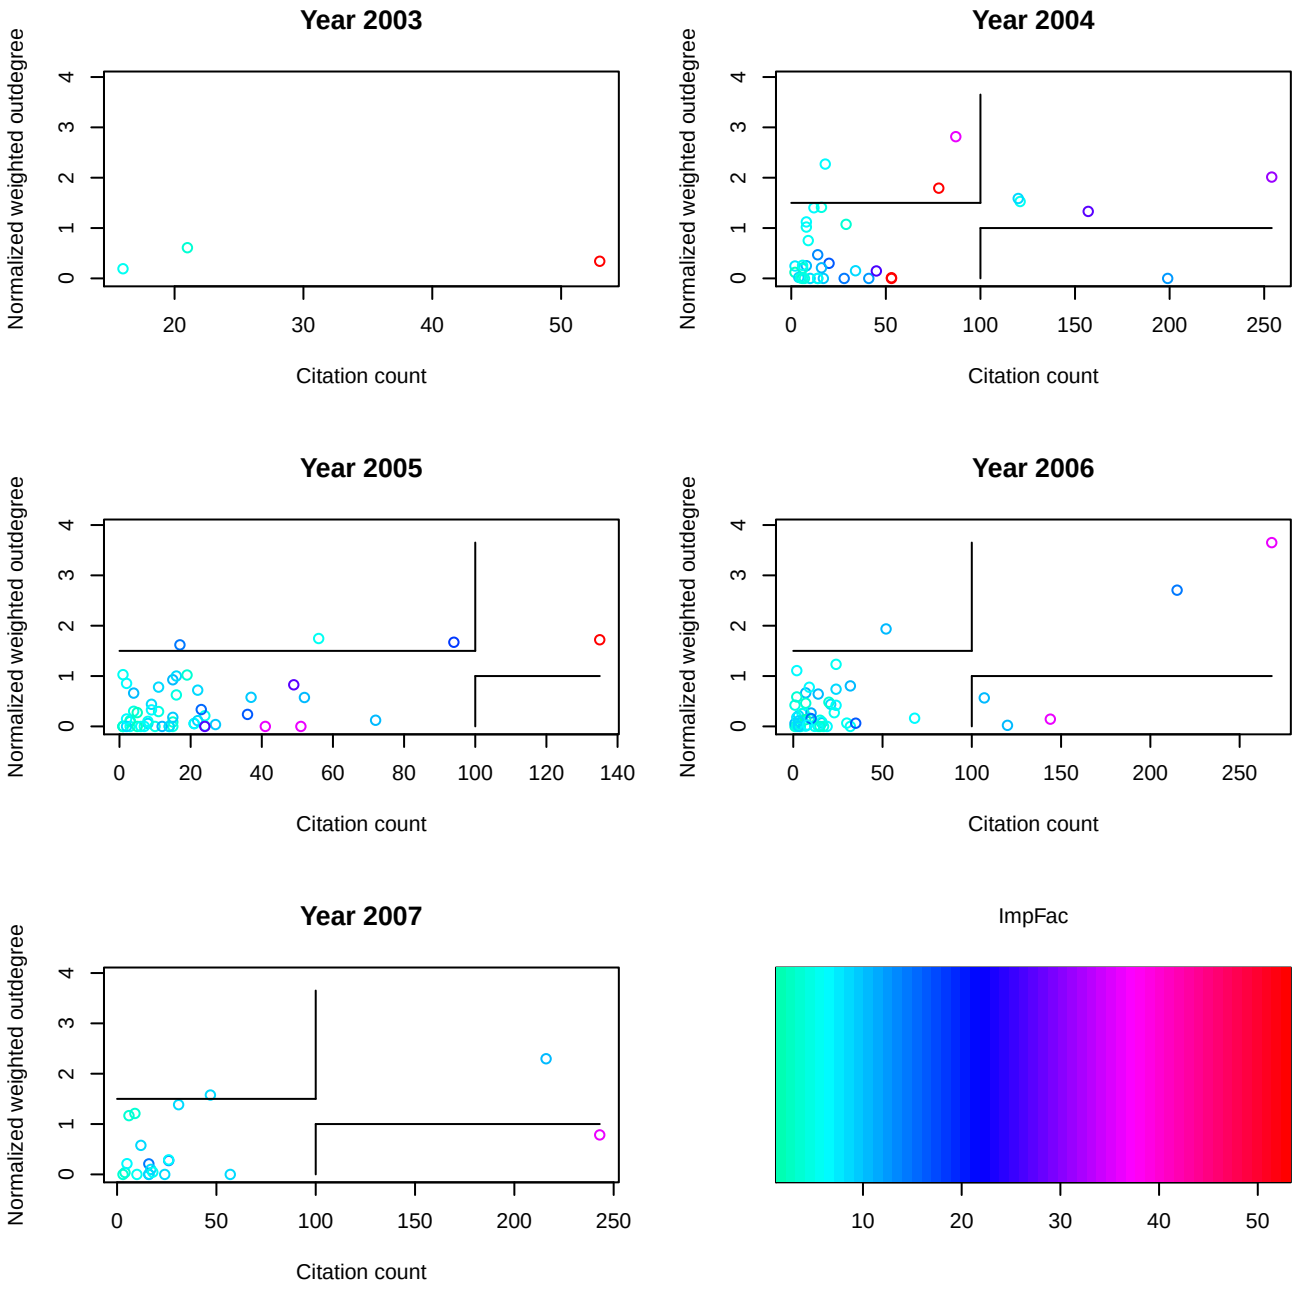

Fig. S1: Stratified data-driven prediction of usefulness of datasets vs. their citation counts. Black solid lines mark the boundary for potentially interesting datasets; the boundaries are set to hold the same percentiles of data as in Fig. 3 in the main paper. *ImpFac* stands for Impact Factor of the publication venue.

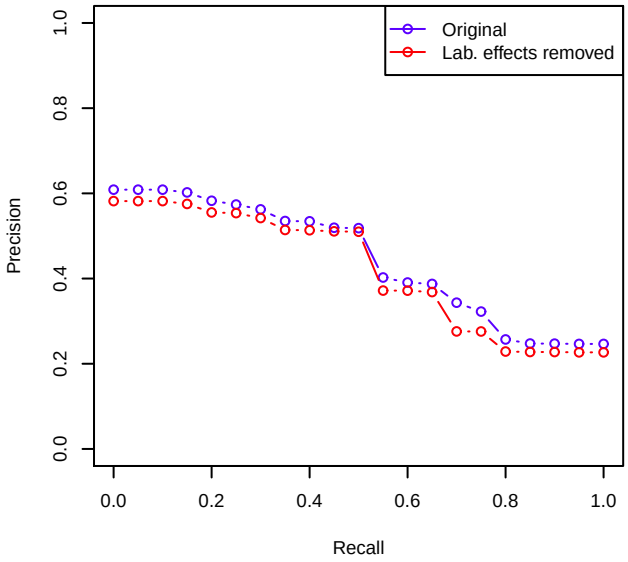

Fig. S2: Removal of laboratory effects changes the retrieval performance only slightly, as measured by the precision-recall curves. *Original:* Replicated from Fig. 1 of the main paper; *Lab. effects removed:* all retrieval results from the same laboratory as the query data have been discarded.

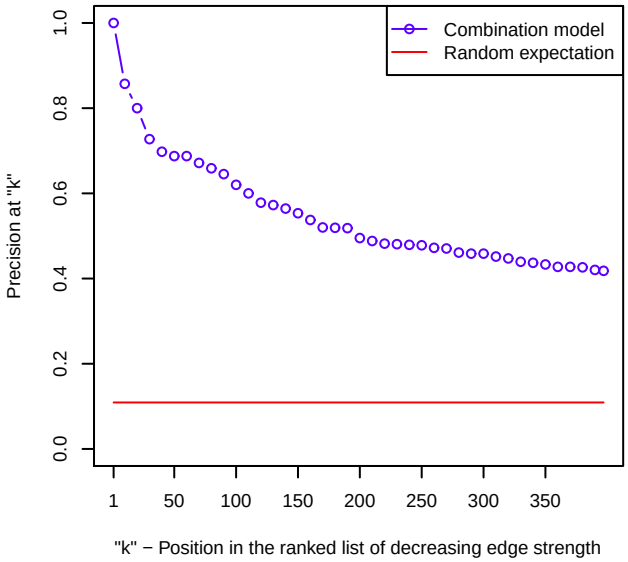

Fig. S3: Overlap of data-driven recommendations with the actual citation graph: Precision *@k* for top edges that explain more than 2.5% variation. The gold standard is the extended citation graph which is built as the union of edges from 1) the original directed graph, 2) between any two articles that are cited together by some other article and 3) between any two articles that have at least one common reference.

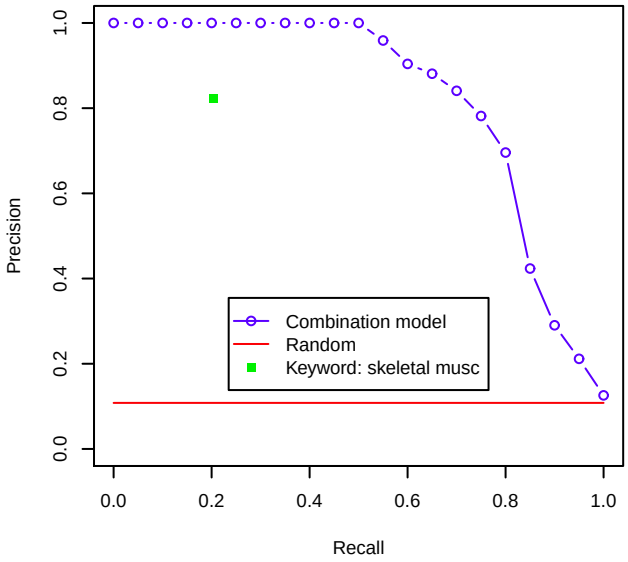

Fig. S4: Retrieval performance evaluation of the data-driven model against keyword search in the skeletal muscle case study. The precision-recall curves are averaged across the 16 skeletal muscle datasets having at least 10 samples.

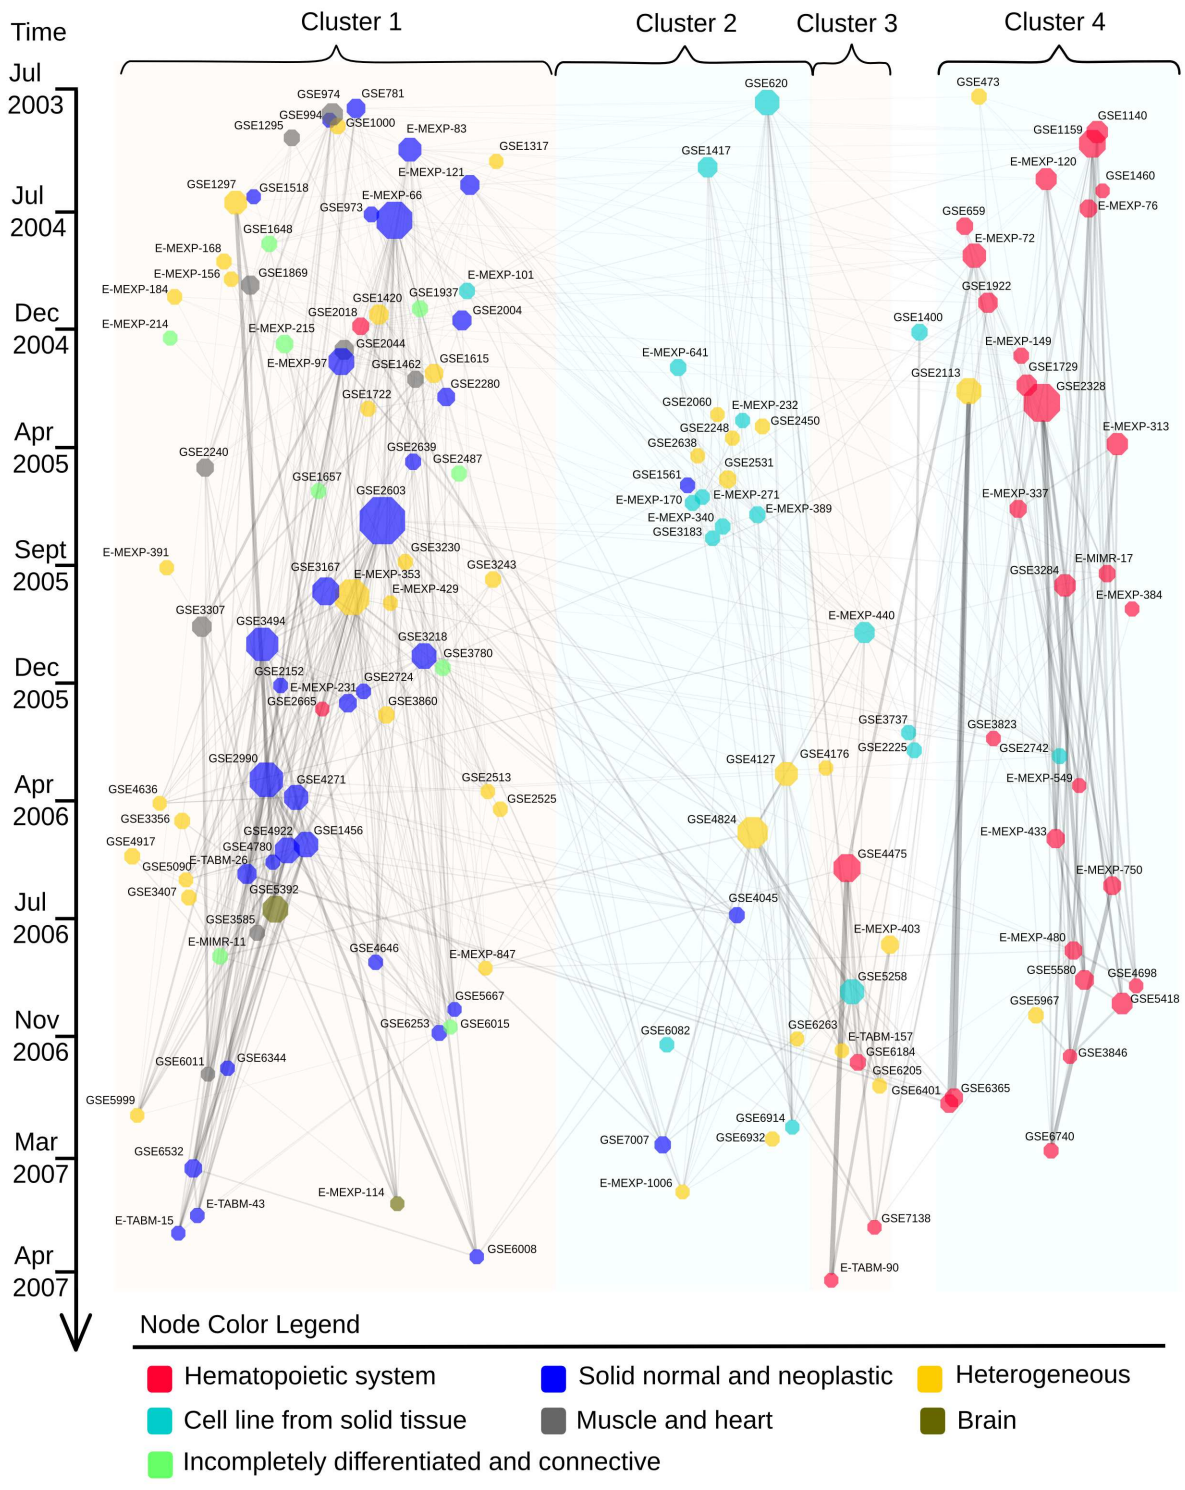

Fig. S5: Enlarged visualization of the relevance network of datasets in the human gene expression atlas (Fig. 2 in the main paper); each dataset was used as a query to retrieve earlier datasets. A link from an earlier dataset to a later one means the earlier dataset is relevant as a partial model of activity in the later dataset. The link width is proportional to the normalized relevance weight (combination weight θ_j^q; only links with $\theta_j^q \geq 0.025$ are shown, and datasets without links have been discarded). The node size is proportional to the estimated influence, *i.e.*, the total outgoing weight. Colors: tissue types (the six meta tissue types [15]). The node layout was computed from the data-driven network using the cluster-emphasizing Sammon mapping scheme. Cytoscape [16] was used to generate the figure.

Table S1: Top 20 strongest edges in the relevance network.

Edge	Clique membership
GSE4475 - E-TABM-90	Does not form a clique
GSE2113 - GSE6365	Does not form a clique
GSE2113 - GSE6401	Does not form a clique
E-MEXP-750 - GSE6740	Part of the CD4+ T helper cell clique
GSE1297 - GSE5392	Part of the brain tissue clique
GSE2328 - E-MEXP-433	Extends the leukocyte clique
GSE2990 - GSE6532	Part of the breast cancer clique
GSE2328 - GSE5580	Part of the leukocyte clique
GSE2990 - GSE1456	Part of the breast cancer clique
GSE3494 - GSE2990	Part of the breast tissue clique
GSE2990 - GSE6008	Does not form a clique
GSE4824 - GSE5258	Part of the adenocarcinoma clique
E-MEXP-72 - GSE4475	Does not form a clique
E-MEXP-403 - E-TABM-90	Does not form a clique
GSE3494 - GSE1456	Part of the breast tissue clique
E-MEXP-337 - E-MEXP-549	Part of the developmental stage of lymphocytes clique
GSE3284 - E-MEXP-433	Part of the leukocyte clique
GSE3307 - GSE6011	Part of the skeletal muscle datasets clique
GSE2990 - GSE4922	Part of the breast cancer clique
GSE2044 - GSE3307	Part of the skeletal muscle datasets clique

Table S2: ArrayExpress accession numbers of 16 skeletal muscle datasets used in the retrieval case study in addition to the human gene expression atlas [15]. All datasets were measured with the human genome platform HG-U133A, the same used in the atlas.

E-CBIL-30	E-GEOD-10760	E-GEOD-11686	E-GEOD-11971
E-GEOD-12648	E-GEOD-1295	E-GEOD-1786	E-GEOD-3307
E-GEOD-4667	E-GEOD-474	E-GEOD-6011	E-GEOD-7146
E-GEOD-8441	E-GEOD-9105	E-GEOD-9397	E-GEOD-9676

Table S3: Skeletal muscle queries with at least one retrieved non-skeletal muscle dataset, sorted according to decreasing precision.

Query	No. of retrieved datasets	No. of retrieved skeletal muscle datasets	Top false positive dataset (rank)
E-GEOD-4667	20	18	GSE5392 (15)
E-GEOD-11686	26	22	GSE5392 (22)
E-GEOD-7146	25	21	GSE781 (22)
E-GEOD-11971	25	19	GSE5392 (15)